POLICY RESEARCH WORKING PAPER   4816

The Developing World's Bulging (but Vulnerable) "Middle Class"

Martin Ravallion

Abstract

The "developing world's middle class" is defined here as those who are not poor when judged by the median poverty line of developing countries, but are still poor by US standards. The "Western middle class" is defined as those who are not poor by US standards. Although barely 80 million people in the developing world entered the Western middle class over 1990-2002, economic growth and distributional shifts allowed an extra 1.2 billion people to join the developing world's middle class. Four-fifths came from Asia, and half from China. Most of the new entrants remained fairly close to poverty, with incomes now bunched up just above $2 a day. The vulnerability of this new middle class to aggregate economic contractions is evident in the fact that one in six people in the developing world live between $2 and $3 per day. Over time, the developing world has become more sharply divided between countries with a large middle class and those with a relatively small one, with Africa prominent in the latter group. Poor people in countries with smaller middle classes may well be more exposed to slowing economic growth.

The Developing World's Bulging (but Vulnerable) "Middle Class"

Martin Ravallion[1]

Development Research Group, World Bank
1818 H Street NW, Washington DC, 20433, USA

Introduction

As the financial crisis emanating from the US spills over to the developing world in 2008/09, many people are asking whether there will be a reversal of the recent progress against poverty. This will depend in part on the distribution of the impacts of the crisis, and it should not be presumed that the poorest will be affected most; ironically, the same things that have kept many people poor in the first place—geographic isolation and poor connectivity with national and global markets—will help protect them from this crisis.[2]

The prospects for reversing past progress will also depend on the distribution of past gains, in so far as this will determine how many people might be vulnerable to even small income losses. Consider Figure 1, which plots three hypothetical cumulative distribution functions (CDFs). The initial distribution is marked "A". If all incomes increase by a similar proportion then the distribution shifts to "B". There will be a fall in poverty across all possible poverty lines and a wide range of measures.[3]

Alternatively (and there are other alternatives), the gains may be larger at lower poverty lines, with little or no reduction in poverty at high lines, as in distribution "C." Poverty is also unambiguously lower for C than A, but C has a bigger bulge in the middle of the distribution, in that the density is appreciably higher at the median than for B and this holds in a wide interval around the median. In terms of the prospects of falling below relatively low poverty lines, distribution C has more people vulnerable to an aggregate economic contraction than B.

So it is important to know whether the developing world has moved toward a distribution more like C than B. While there have been a number of empirical studies of global poverty and inequality,[4] this issue has been neglected.

This paper aims to fill this gap in knowledge. The following section discusses the generic issues that arise in measuring the size of the "middle class" in developing countries and proposes a new definition. The paper then presents empirical results on the size of the middle class, how it has changed over time, and the proximate causes of the observed cross-country differences in the pace at which the middle class has grown. Some observations are also made on the possible

---

[2]   For evidence on this point in the context of the longer-term impacts of the Indonesian financial crisis in 1998 see Ravallion and Lokshin (2007).
[3]   As shown by Atkinson (1987), if the CDFs for two distributions do not intersect up to some maximum poverty line, then the ordinal comparison is robust to both the poverty line and the precise poverty measure, within a broad class of additive measures.
[4]   Contributions include Bourguignon and Morrison (2002), Chen and Ravallion (2004, 2008), Milanovic (2005) and Sala-i-Martin (2006); for a recent survey of this literature see Ferreira and Ravallion (2009).

implications of these empirical findings for assessing the poverty impacts of the diminished growth prospects of developing countries, in the wake of the 2008/09 crisis.

1. Measuring the middle class

There is a large literature on defining and measuring the "middle class" in the context of developed countries, where the main issue has been the (claimed) decline in the middle class in the US (and some other Western countries) over recent decades. Typically, being "middle class" is defined as having an income within some interval that includes the median and the interval has typically been symmetric in the income space around the median. The lower and upper bounds have been set in diverse and *ad hoc* ways. An influential early paper by Thurow (1987) focused on the interval from 75% to 125% of the median. The literature has shown signs of converging on Thurow's definition; for example, this is Pressman's (2007) definition in his study of whether there has been a decline in the middle class for 11 developed countries, including the US.

Amongst the studies that have used data for developing countries (sometimes with developed countries), the main issue has been the role of the middle class in promoting economic development, such as through entrepreneurship, policy reform and institutional change.[5] But there is little sign of agreement on what the "middle-class" means. In keeping with much of the literature on the US, Birdsall et al. (2000) defined the middle class as those with incomes between 75% and 125% of the median in each country. In contrast to this relative definition (with real-income bounds specific to each country), other authors have defined the middle class in purely absolute terms (with common bounds across countries). Milanovic and Yitzhaki (2002) defined the middle class as those living between the mean incomes of Brazil and Italy,[6] while Banerjee and Duflo (2008) identified the "middle class" in developing countries as those living between $2 and $10 a day. The latter two studies do not even use overlapping intervals; nobody is likely to be "middle class" by both the Milanovic-Yitzhaki and Banerjee-Duflo definitions.[7]

In large part this is a matter of whether one is after a definition appropriate to rich countries or poor ones. People living below $10 a day would clearly not be considered "middle class" in most developed countries; indeed, they would be living well below the US poverty line,

---

[5] See, for example, the discussions in Birdsall et al. (2000), Easterly (2001), Sridharan (2004) and Banerjee and Duflo (2008).

[6] Also see Bussolo et al. (2008) (and their results reported in World Bank, 2007) who use the Milanovic-Yitzhaki definition in identifying a "global middle class."

[7] Brazil's mean consumption per capita (from national accounts) was about $12 per day in 2005.

which was about $13 a day in 2005.[8] Yet it is likely that many people in developing countries living below the US poverty line, and the Milanovic-Yitzhaki standard, would be deemed "middle class". Consider the two most populous countries. The closest concept to "middle class" in China is "Xiaokang;" eventually achieving the "Xiaokang society" is the goal of China's reforms, instigated in 1979 under Deng Xiaoping's leadership. The Government of China's National Bureau of Statistics set a minimum income for Xiaokang in 1991; when converted to 2005 PPP $'s, Xiaokang requires $2.24 per day in rural areas and $3.47 a day in urban areas; I estimate that over 500 million Chinese were Xiaokang by 2005 (using *PovcalNet*), which (as we will see) far exceeds the number living above the US poverty line. It is clear that many people who would be deemed "poor" in the US are thought of as middle class in China.

This is true in India too. It is often claimed that 300 million people are now "middle class" in India; see, for example, the Wikipedia entry on the "Standard of Living in India" (although I have had little success in tracking down the origin of that number). The surveys done by the National Council of Applied Economic Research (NCAER) are often used in defining India's middle class; based on that source, Sridharan (2004) gives a range of definitions implying that 100-250 million people are middle class in India around 2000. From the most recent NCAER survey, Shukla (2008) deems 25 million households (about 120 million people) to be middle class in 2007-08. As we will see, all these estimates far exceed the likely number of people in India who are not poor by US standards.

One can also question the relevance of other definitions found in the literature. It seems implausible that a definition of the middle class relevant to developing countries would be centered on the median, which might more reasonably be deemed a lower bound. And, while some degree of arbitrariness in such measurement efforts is inevitable, one might wonder what meaning can be attached to the Banerjee-Duflo bounds of $2 and $10.

The definition proposed here is close to the Banerjee-Duflo definition, but I try to give the bounds more meaning. The developing world's middle class is defined as those who are not deemed "poor" by the standards of developing countries but are still poor by the standards of rich countries. For the lower bound I will use the median poverty line in poverty lines for 70 national poverty lines, drawn from in-country poverty measurement studies by the World Bank and national governments; the data on national poverty lines are described in Ravallion et al. (2008).

---

[8]   I have used the line for a family of four from the web site of the US Department of Health and Human Services.

Each of the national poverty lines in this sample is designed to attain recommended food-energy requirements with (socially-specific) allowances for basic nonfood needs. The median of these national lines is $2.00 per day at 2005 purchasing power parity (PPP) using the PPPs for household consumption from the 2005 International Comparison Program (World Bank, 2008; www.worldbank.org/data/icp). As we will see, this is only slightly below the median for the population of the developing world as a whole.

By definition, half the developing countries in the Ravallion et al. sample have national poverty lines below $2 a day.[9] For example, the official poverty lines in both China and India are only about half this figure, so there are many people who are not poor by official standards in both countries who are not yet "middle class" by my definition. Nonetheless, there can be no doubt that $2 a day is a frugal lower bound even amongst developing countries. There is a marked economic gradient in national poverty lines, reflecting more generous (food and nonfood) allowances for "basic needs" as countries develop (Ravallion et al., 2008). The highest line in the Ravallion et al. sample of developing countries is $9 a day (for Uruguay). For some purposes, I will also focus on what might be termed the developing world's upper middle class, for which the lower bound is $9 a day.

As an upper bound for defining the developing world's middle class, I will use the aforementioned US poverty line of $13 a day at 2005 PPP. This is a long way from the reach of those who are poor by developing-world standards. At an income growth rate of (say) 5% per year it would take almost 40 years to bring someone living on $2 a day up to the US poverty line; even at a growth rate of 10% it would take 20 years. Clearly $2-$13 is a wide interval.

I will not say that someone has entered the Western middle class until the person has reached the US poverty line, which accords closely with the lower bound of the Milanovic-Yitzhaki definition.

When talking about a growing economy, finding more density toward the middle could simply be due to that growth. It is also of interest to isolate the extent of the middle-income bulge that is attributable to changing distribution. There are a number of ways one might do this. The literature on the rise of the middle class in the US has often used a standard inequality measure. However, such a measure need not reflect well the extent to which density has shifted

---

[9] For measuring extreme poverty, Ravallion et al. (2008) recommend a lower line of $1.25 a day. This is the average poverty line of the 15 poorest countries in the world.

relative to the median; in particular, Wolfson (1994, 1997) points out that shifts in density away from the median, and from the lower and upper extremes, can be generated by transfers that entail <u>lower</u> inequality (by the usual transfer axiom used in inequality measurement).

A better measure from this point of view is the Foster-Wolfson (FW) "polarization index," based on the spread relative to the median of incomes normalized by the median.[10] More precisely the polarization index is the area under the "polarization curve" given by the vertical distance between the Lorenz curve and the tangent to that curve at the median. This gives how far each percentile is from the median. The emergence of a middle-income bulge due to a change in distribution can be interpreted as a downward shift in the polarization curve. To avoid confusion about "which way is up" I will use 100 minus the FW index (in %), and I shall refer to this as the FW index of concentration around the median rather than "polarization." This provides a convenient summary statistic, although (as noted above), it is questionable whether the median is the relevant anchor in poor countries.

However, there is a more natural, and less restrictive, way of assessing the contribution of distributional changes to the "bulging middle," namely to construct a counterfactual for the second date in which the Lorenz curve does not change relative to the base date, but the overall growth rate is the same as that observed in the data.[11] I will also examine whether there is a bulging middle relative to this distribution-neutral counterfactual.

2.   Evidence on the developing world's expanding middle class

Using almost 700 household surveys for 115 developing countries, Chen and Ravallion (2008) estimate poverty measures for a range of international poverty lines anchored to the definitions of "poverty" found in practice in poor countries. The salient features of the Chen-Ravallion estimation methods are as follows. Each international poverty line was converted to local currencies in the benchmark year (2005) using the consumption PPP for 2005 from the ICP. It was then converted to the prices prevailing at the time of the relevant household survey using the best available consumer price index for that country. Then the poverty rate was calculated from that survey by standard methods. Interpolation and extrapolation methods were used to line

---

[10] Wolfson (1994) refers to an earlier, un-published, paper with James Foster as the source of the concepts of a polarization curve and polarization index.
[11] This approach is in the spirit of the Datt and Ravallion (1992) decomposition of changes in poverty measures into "growth" and "redistribution" components.

up the survey-based estimates with common reference years across countries, including 2005. Population-weighted aggregate measures were then formed by region and globally.

Table 1 reproduces the Chen-Ravallion estimates of the poverty rate for the $2 a day line. Some notation will help for brevity. Let $F_t(z)$ is the poverty rate (often called the "headcount index") at date t=(19)90 and (20)05 for the poverty line z in $'s per day at 2005 PPP. So Table 1 gives $F_{90}(2)$ and $F_{05}(2)$ by region and in the aggregate for the developing world.

On repeating these calculations using the $13 a day line, I find that over 95% of the developing world in 2005 is deemed to be poor by US standards; Table 2 gives my estimates of $F_{90}(13)$ and $F_{05}(13)$. It is striking how few people in the developing world are not poor by US standards. Possibly most notable is the count for India, which implies that only three million people are not poor by this standard. Under-reporting of consumptions in India's National Sample Surveys, or selective compliance in the survey's random samples, could well be leading to an under-counting of those not poor by US standards. But even if we use a poverty line of $9 a day—equivalent to using the US poverty line but increasing all consumption levels by about 50%, which would seem a very large adjustment for measurement error—the count of those in India who are not poor by US standards rises to only six million in 2005.

It is no less striking how little $F_t(13)$ has fallen very little over time—by only one half of a percentage point (Table 2). So the number of people living in poverty by US standards has risen over the period. In 2005, there were only 250 million people in the developing world who were not poor by US standards, rising from 170 million in 1990. Defining the "Western middle class" as those living above the US poverty line, the developing world added 80 million to the count of the middle class by this definition over 1990-2005.

Notice that, on comparing Tables 1 and 2, China's poverty ranking changes relative to the rest of the developing world. Using the $2 line, China had a lower incidence of poverty than the average excluding China in 2005, but this reverses using the $13 line; less than 2% (25 million people) of China's population had attained the US poverty line in 2005.[12] China's accounted for about one quarter of the increase in the size of the Western middle class over this period.

---

[12] The 2005 ICP round was the first time that China had participated officially in the ICP, and the price data imply a higher PPP than past estimates from other sources, implying that China is poorer than we thought. The estimates reported here have attempted to correct for the "urban bias" in the sample frame of the 2005 ICP for China; for details see Chen and Ravallion (2009).

It is plain from Tables 1 and 2 that a large middle-income bulge emerged. Table 3 gives the population share of the developing world's middle class—the proportion living between $2 and $13 a day ($F_t(13) - F_t(2)$). This increased from 33% in 1990 to 49% in 2005, or from 1.4 billion to 2.6 billion people. China accounts for half of this increase and East Asia as a whole for two-thirds. In 1990, only 15% of the Chinese population lived above $2 but below $13; by 2005 the proportion had risen to 62%. Though not as large (absolutely or proportionately) as for China, India saw an extra 117 million people join this middle-income group. Asia as a whole (east and south) accounts for 81% of the expansion in the number living between $2 and $13 per day. All regions of the developing world saw an increase in $F_t(13) - F_t(2)$ except Eastern Europe and Central Asia where the number fell slightly over the period.

Where is the bulge within the (2, 13) interval of incomes? The empirical estimates of the CDF's and densities for 1990 and 2005 are given in Figure 2.[13] While a reduction in poverty is indicated for all lines, it is clear that this has not been a simple rightward displacement of the density function. Indeed, the mode has remained almost unchanged over this period, at around $1.00 per day. The mean and median have increased, from $3.14 and $3.94 per day for 1990 and 2005 respectively for the mean, while the medians were $1.47 and $2.13. Instead of a simple rightward displacement, we have seen a marked "bunching up" due to a shift in densities from below the $2 a day line to just above it, with the bulk of the gain in the interval $2 to $6 a day. The bulge in 2005 is mostly above the median. Figure 3 gives $F_{05}(z) - F_{90}(z)$ for $0 < z \leq 13$, i.e., the extent to which the poverty rate has fallen according to different poverty lines up to the US line. It can be seen that the reduction in the poverty rate peaks at about $1.50 a day (almost exactly the 1990 median). The impact on the poverty rate falls below 5% points at poverty lines of about $6 a day or higher.

It is clear that very little of the expansion in the developing world's middle class was due to its "upper middle class," namely those living above the highest poverty line found amongst developing countries, but still below the US line. Over 1990-2005, the estimated value of $F_t(13) - F_t(9)$ rose from 3.1% to 4.3%, or from 139 million to 233 million. Of the extra 1.2 billion people who joined the middle class, only 95 million made it to this upper stratum.

---

[13] The densities were estimated by taking the empirical derivatives of the estimated CDFs derived using *PovcalNet* for each of 25 selected poverty lines.

The shift in the density functions reflects in part the overall positive growth in the mean. Recall that two methods are to be used for isolating the contribution of distributional shifts. The first method is the Foster-Wolfson index of concentration around the median. This rose from 61.3% to 63.9% over 1990-2005, indicating that the changes in the distribution of incomes normalized by the median entailed higher concentration around the median.

The second method uses the aforementioned distribution-neutral counterfactual, as in Figure 3. Here I have calculated the poverty impacts when the distribution moves horizontally according to the proportionate increase in the mean between 1990 and 2005, i.e., all 1990 income levels are scaled up by the same growth rate, leaving relative distribution at its 1990 level.

Comparing the two curves in Figure 3, it is evident that the actual changes in relative distribution had substantial impacts on poverty, as judged by developing country standards. Under the counterfactual we would have seen lower poverty impacts at low poverty lines and higher impacts at high lines. When assessed relative to this counterfactual, the bulge now starts below the 2005 median and persists until a turning point at about $8 per day. If we define the bulge as more than a 2 percentage point gap between the actual and counterfactual distributions then it spans the interval $1.00-$5.00 per day, containing 63% of the developing world's population. (The region $1.25-$5.00 contains 53% of the population.) If instead we define the bulge as more than a 6 percentage point gap then the region narrows to about $1.50-$3.00 a day (or about 0.6 median to 1.4 median); 30% of the population is in this interval.

Focusing on the $2 a day line, the actual growth process of the developing world implied a 25.8% (proportionate) reduction in the poverty rate over 1990-2005 (Table 1). The implied elasticity of poverty reduction with respect to the overall survey mean is -1.0. By contrast, the counterfactual growth process implies a 15.0% drop in the $2 a day poverty rate and an elasticity of -0.6. The overall growth process in the developing world has clearly been more pro-poor than implied by distribution-neutral growth.

The elasticity varies markedly with the poverty line. Figure 4 gives the elasticity of $F_t(z)$ with respect to the mean holding relative distribution constant across the full range of poverty lines; given that $F_t(z)$ is homogeneous of degree zero in the mean and z the point elasticity is given by $-zf_t(z)/F_t(z)$ (Kakwani, 1993). The absolute elasticity falls sharply at poverty liens under $2 a day, from 3.0 at a line around $1.00 a day to 1.0 at around $2. The figure also gives the elasticities that would have been obtained in 2005 if the growth process over 1990-2005 had

9

been distribution-neutral for the developing world as a whole. The distributional shifts have raised the elasticity across all line, but only noticeably so over about $2 a day.

3.   Bulging middles <u>within</u> developing countries?

For each of almost 100 developing countries, it is possible to assess whether there has been a bulging middle over time. I have constructed the longest spells between two surveys, both using the same indicator, either consumption or income per person.[14] The median year of the first survey is 1991 while the median for the second is 2004. The median interval between surveys is 13 years and the interval varies from three to 27 years. All changes between the surveys are annualized. Again I use both the changes in the population share between $2 and $13 a day and the changes in the FW index of concentration around the median. Recall that the former measure directly reflects growth in the mean, while the latter is a distributional measure; I also examine distributional changes in the former measure. Precise sample sizes vary depending on the variables used in the analysis.

Some further notation will be useful. Let the absolute and proportionate annualized differences for the variable x be $d_i(x_{it}) \equiv (x_{it} - x_{it-\tau_i})/\tau_i$ and $g_i(x_{it}) \equiv (\ln x_{it} - \ln x_{it-\tau_i})/\tau_i$ respectively, for country i looking back from date t over a time interval of length τ_i. Note that $g_i(.)$ is simply the average growth rate of x over the period, so $g_i(\mu_{it})$ is the rate of growth in the mean, μ_{it}. Also let $F_{it}(z)$ denote the CDF for country i at date t, giving the proportion of the population living in households with consumption or income not exceeding z in $'s per day at 2005 PPP. (So $F_t(z)$ is the population weighted mean of $F_{it}(z)$ over all i.) In measuring the size of the middle class in a given country, I will use both the absolute difference $A_{it} \equiv F_{it}(13) - F_{it}(2)$ and the relative difference, $R_{it} \equiv F_{it}(13)/F_{it}(2)$. Note that $g_i(R_{it})$ is also the difference in growth rates for the poverty measures evaluated at $13 and $2 a day.

Sizes of the middle class across countries

Table 4 gives a 5x5 contingency table that summarizes the joint distribution across the two survey rounds of the sizes of the middle class across countries. A strong positive correlation

---

[14] When both were available, consumption was generally preferred The only exception was Peru, for which a much longer time period could be covered using income-based distributions.

is evident.[15] There are two main off-diagonal elements. The first is the expanding middle-class amongst those countries that started with middle-class shares less than 40%. Secondly there is a marked fall in the size of the middle-class in some countries that started with a large share (over 80%) of the population in the $2-$13 interval; some of this is undoubtedly measurement error.

Figure 5 plots the size of the middle class for the latest survey against the earliest (i.e., A_{it} against $A_{it-\tau}$). The regression line has a slope significantly less than unity; the regression coefficient is 0.704 with a robust standard error of 0.050, and the intercept is 18.776% (s.e.=3.149).[16] (The fixed point solution is at a middle-class share of 63.5%.) The share of the population living between $2 and $13 a day rose in 64 out of the 99 countries. The mean $d_i(A_{it})$ was 0.22 percentage points per year and the median was 0.45 percentage points per year. The mean $g_i(R_{it})$ is 0.024 (and the median is 0.011) or a growth rate of about 2% per year.

There is a markedly bimodal distribution across countries in the population shares living between $2 and $13 a day; this is evident in the clustering of data points in Figure 5 but is easier to see in the kernel density functions for A_{it} and $A_{it-\tau}$ in Figure 6.[17] Taking a share of 40% as the cut-off point, 30 countries are in the lower mode and 69 are in the upper one for the most recent survey; the corresponding counts for the earliest surveys are 42 and 57 (Table 4). Over time, there has been a shift of density toward the upper mode, away from the lower one.

There is a notable regional concentration of Sub-Saharan African (SSA) countries in the lower mode. Two-thirds (20 out of 30) of the lower-mode countries in Figure 6 for the most recent survey rounds in SSA (versus 29% for the sample as a whole) and over two-thirds (20 out of 29) of the SSA countries in the sample are in the lower mode. In the latest survey rounds, the mean population share in the $2-$13 interval for SSA was 27.6% (s.e.=3.5%) as compared to 63.5% (s.e.=2.3%) in the set of all other countries. In the earliest survey rounds, the corresponding means were 22.9% (4.3%) and 58.8% (2.9%).[18] (In both cases, the difference between SSA and non-SSA countries is statistically significant at the 1% level.)

---

[15] Cramer's V statistic is 0.453 and the contingency coefficient is 0.671. Pearson's test gives a Chi-square of 81.35, which is significant at better than the 1% level.

[16] All standard errors reported in this paper are corrected for heteroskedasticity using White's method.

[17] The fact that Figure 4 suggests "convergence" to a common signal middle-class share yet there is marked bimodality in the densities echoes Quah's (1996) critique of past interpretations of "growth regressions."

[18] The population share in the $2-$13 interval is also correlated with whether the survey used consumption or income, with lower share for consumption surveys. However, this difference vanishes if one controls for SSA countries; the correlation stems from the fact that SSA countries have mostly used consumption.

Does the middle class expand or contract as an economy grows? Plainly that depends on whether growth is more effective in reducing poverty by the standards of poor countries, or poverty by US standards. Empirically, the answer is clear: the pace at which the middle-class expands tends to be higher in more rapidly growing economies. The correlations are higher for the proportionate changes in the middle's population share; the annualized change in the log of the ratio of the proportion living below \$13 to that below \$2 has correlation coefficients of 0.55 and 0.37 with the growth rates in the survey means and national accounts (NAS) consumption per capita respectively. The corresponding correlations with the annualized changes in the levels of the population shares living between \$2 and \$13 are 0.28 and 0.19. Figure 7 plots the proportionate middle-income bulges ($g_i(R_{it})$) against the growth rates in the means, $g_i(\mu_{it})$. The regression coefficient in Figure 7 is 1.248 with a robust standard error of 0.212 (n=91) and R^2=0.333. Evidently this regression coefficient is the difference between the regression coefficient of $g_i(F_{it}(13))$ on $g_i(\mu_{it})$ and that for $g_i(F_{it}(2))$ on $g_i(\mu_{it})$. Thus the positive growth effect in Figure 7 can be interpreted as saying that the growth elasticity of poverty reduction is greater (in absolute value) for the lower poverty line. The regression coefficient of $g_i(F_{it}(2))$ on $g_i(\mu_{it})$ is -1.372 (s.e.=0.231; n=91) while the corresponding regression coefficient for $g_i(F_{it}(13))$ is -0.124 (s.e.=0.036), although it is higher if one uses all the observations (including those for which $g_i(F_{it}(2))$ cannot be calculated given that one or both poverty rates are zero); for that sample the coefficient is -0.20 (s.e.=0.05; n=99).

There is no sign at country level of a non-zero mean distributional effect on the size of the middle-class. This is evident in Figure 7, on noting that the regression line passes though the (0,0) point, implying that the ratio of the population share under \$13 to that under \$2 is roughly stationary at zero growth. (The intercept of the regression in Figure 7 is 0.003 with a standard error of 0.006.) So the overall distributional effect on the expansion in the middle class for the developing world as a whole is driven largely by the between-country effect rather than within-country. Even so, changes in distribution (in both directions) within countries have contributed to the changes in the size of their middle class. This is suggested by the fact that the growth rates in the mean only explains one third of the variance in Figure 7, although measurement errors are likely to be playing a role and it is difficult to distinguish these from idiosyncratic distributional shifts. It is at least suggestive that if one adds to the regression in Figure 7 a control for changes

in inequality (measures by the log difference in the Gini index) this is significant (at the 5% level), with a negative coefficient, and the share of the variance explained rises to one half.

Concentration around the median

The mean Foster-Wolfson index rose only slightly over the period, at 63.9% and 64.9% for 1990 and 2005 respectively; the median increased from 64.6% and 65.6%.[19] The density of the annualized growth rates in the FW index is centered at a small negative value; the mean change is -0.13% per year. The index has fallen slightly more often than it has risen; 48 out of 91 countries show a decrease. The majority of developing countries have not experienced higher density in a neighborhood of the (country-specific) median, although there has been a clear tendency for bunching up above the $2 line. Consider, for example, the two most populous countries, China and India. In marked contrast to the findings in Table 3, China's FW index has tended to fall over time, while India's has remained largely unchanged.[20]

As an aside, it can be noted that, despite the conceptual differences noted by Wolfson (1994, 1997), the FW index is highly correlated with a standard inequality index.[21] The correlation coefficients between the FW index of polarization and the Gini index of inequality in the most recent surveys is 0.94 (and 0.92 for the earliest survey). The correlation coefficient between the changes over time in the two measures is 0.83. However, the Gini index is the better predictor of changes in the poverty rate; if one includes both the change in the log Gini index and the change in the log FW polarization index as extra controls in the regression of $g_i(R_{it})$ on $g_i(\mu_{it})$ (Figure 7) only the Gini index is significant at the 5% level.

There is little sign that changes in the FW index of concentration around the median have been correlated with rates of economic growth over the same time period. The correlation coefficients between the annualized growth rate in the index and the real annualized growth rate in the survey mean is 0.09 while the correlation with the corresponding growth rates in real consumption per capita from the national accounts is 0.00. It is not the case that more rapidly

---

[19] Given that these mean FW indices across countries do no include the between-country component of the distributional change, they are higher than the overall FW indices reported earlier in this paper. Recall also that this is 100 minus the usual FW "polarization" index.

[20] For the longest time periods available, China's FW index fell from 74.5% to 61.1% over 1981-2005, while India's remained fairly similar, at 75.0% and 74.5% for 1977/78 and 2004/05 respectively.

[21] A similar point is made by Zhang and Kanbur (2001), using data for China.

growing developing economies have experienced higher (or lower) rates of increase in the extent of polarization by this measure.

However, there is a strong positive correlation between the initial level of mean consumption and the subsequent change in the FW index; Figure 8 plots the annualized proportionate change in the index against the log of the initial mean.[22] The correlation is just as strong with the initial $2 a day poverty rate (r=-0.33), and the turning point is at a poverty rate of 24.83% (s.e.=8.78%). So there is evidence of a middle-income bulge relative to the median emerging in the poorer countries within the developing world.

Table 5 brings the above observations together in the form of regressions for the growth rates (annualized log differences) in the two measures used above; the regressors are the growth rate and the initial mean, using both survey means and household consumption per capita from the NAS. The growth effects on the bulge in the population shares in the ($2, $13) interval are evident for both surveys means and NAS data, though smaller for the latter. Note that the effect of the initial mean has opposite effects on the two dependent variables. Given that the regression includes the growth rate, the coefficient on the initial mean can be interpreted as a distributional effect. So the finding that the distributional effects in poorer countries tended to generate greater bulge around the median is not robust to the choice of that measure in that focusing on the income space above $2 a day gives the opposite conclusion. This sensitivity to the choice of measure is probably not too surprising given that the FW index is centered on the median while the bulk of the shift in density relative to the distribution-neutral counterfactual is to a region somewhat above the median. In fact, on controlling for the growth rate, I find no significant partial correlation between the two measures (a t-statistic of 0.78 using the growth rate in the survey mean and 1.34 using the growth rate for NAS consumption).

4. Does a larger middle class help assure pro-poor growth?

Looking back over time, the countries with a larger middle-class tended to attain higher subsequent rates of poverty reduction from a given rate of growth. Regressing the rates of poverty reduction ($g_i(F_{it}(2))$) on both the rate of growth ($g_i(\mu_{it})$) and its interaction with either the absolute share of the population in the middle class ($A_{it-\tau}$) or the log relative share ($\ln R_{it-\tau}$) I

---

[22] Figure 7 uses the survey means. The correlation is also found using private consumption per capita from the national accounts (r=0.31; prob.=0.006). The correlation is also robust to adding a regression control for the initial level of the polarization index.

found that the interaction effect has a significant negative coefficient. However, I also found that this effect was almost entirely attributable to the poverty rate for the $2 line; it is the differences in the extent of poverty by this line that account for the bulk of the cross-country variance in the size of the middle class.

So the (absolute) growth elasticity of poverty reduction tends to be lower in countries with a higher initial poverty rate.[23] Poorer countries tend to experience lower proportionate impacts on their poverty measures from any given rate of growth. Furthermore, the countries with high current poverty rates have tended to experience less middle-income bulge.[24] All this suggests that those countries that have enjoyed the largest expansions in their middle class that will be most vulnerable to aggregate economic downturns in the wake of the crisis.

However, there is an important caveat to this argument. It is not the proportionate impact on poverty that one is most concerned with in this context, but the absolute impact, for that is what matters to the number of people experiencing extreme poverty. When talking about the developing world as a whole, this is not an issue. But it is relevant to cross-country comparisons. Furthermore, the tendency for growth to be less poverty-reducing in poorer countries vanishes when one focuses instead on the absolute impacts. This is evident from the regressions reported in Table 6. The left panel (regressions (1)-(4)) gives regressions of the annualized change in the log of the $2 a day poverty rate against both the annualized growth rate in the mean and its interaction with the initial poverty rate. Both Ordinary Least Squares (OLS) and Instrumental Variables (IV) estimators are given; the IV method uses the growth rate in private consumption per capita from the NAS as the instrument for the growth rate in the survey mean.[25] The table

---

[23] Mathematically, this elasticity will be a function of the mean (relative to the poverty line) and the properties of the Lorenz curve, roughly interpretable as "inequality." Ravallion (1997) argued that in high inequality countries the poor will tend to have a lower share of the gains from growth and showed that this holds empirically; this was substantiated in a larger data set by Ravallion (2007). Ravallion (1997) did not, however, find empirically that the growth elasticity of poverty reduction varied systematically with the mean. Lopez and Serven (2006) show that if incomes are log-normally distributed, then such a variation with the mean is implied theoretically, and they show that log-normality is a good approximation empirically. Easterly (2009) conjectures that the initial poverty rate is likely to be the better predictor of the elasticity though no evidence is provided. The results reported in the present paper, on a larger and better quality data set than used by Ravallion (1997, 2007), suggest that differences in the initial mean are important for explaining differences in the elasticity, and that the initial poverty rate is an even better predictor than the initial mean, consistently with Easterly's conjecture.

[24] The correlation coefficient between the log of the $2 a day poverty rate in the most recent survey and the proportionate increase in the share of the population between $2 and $13 a day is 0.620.

[25] Following Ravallion (2001), this IV allows for the possibility that a spurious negative correlation exists due to common measurement errors (given that the poverty measure and the mean are calculated from the same surveys)

also gives homogeneity tests, which pass comfortably; by implication, the relevant growth rate is the poverty adjusted rate, as given by the growth rate *times* one minus the poverty rate.

The positive interaction effects between the growth rate and the initial poverty rate for the regressions (1)-(4) in Table 6 confirm that the (absolute) growth elasticity of poverty reduction tends to be lower in poorer countries.[26] However, when we turn to the right panel (regressions (5) and (6)) the interaction effect vanishes—indeed it switches sign. In terms of the absolute impacts on poverty, higher growth tends to be more poverty reducing in poorer countries, not less. Conversely, the poorer the country, the greater the absolute impact on poverty that is to be expected from any given rate of aggregate economic contraction.

This distinction between the proportionate and absolute impacts of economic growth on poverty has an important implication for development efforts to attain the first Millennium Development Goal (MDG1), namely to halve the 1990 poverty rate by 2015. Easterly (2009) argues that MDG1 is harder to achieve in Africa because the growth elasticities of poverty reduction are likely to be lower (in absolute value) than elsewhere, given that poverty incidence tends to be higher. The results of the present paper make clear that this argument depends crucially on the fact that MDG1 calls for a common proportionate decrease.

When applied to specific countries, this is an odd feature of MDG1. Aiming for a common proportionate decrease is plainly inconsistent with an aid strategy that focuses on reducing aggregate poverty (no matter where one happens to live) using aid allocations across countries that differ in the relevant parameters (growth prospects, aid effectiveness in promoting growth, growth elasticities of poverty and initial levels of poverty).[27] While it might be a moot point how much aid allocations in practice are guided by MDG1 (or aggregate poverty reduction), it can be argued that focusing on MDG1 at country level is not only unfair to Africa but would also distort aid policies against attaining MDG1 in the aggregate.

---

[26] It might be conjectured that this is an "Africa effect" given past findings in the literature indicating that SSA tends to have lower growth elasticities of poverty reduction (Chen and Ravallion, 2004, 2008). However, the effect persists (with similar size and significance) if one adds a dummy variable for SSA countries (both on its own and interacted with the growth rate) to these regressions.

[27] The marginal impact of external aid (per capita) to a given country on the number of poor in the world is given by the product of the (country-specific) values of the growth elasticity of poverty reduction, the marginal impact on log mean income of extra aid and the initial poverty rate in that country. Minimizing aggregate poverty by equating marginal impacts would clearly entail differing proportionate impacts on poverty, depending on the country-specific parameters.

5.   Conclusions

Like "poverty," the term "middle-class" is defined differently in different countries at different levels of economic development. Some observers have applied a rich-world concept of what it means to be middle class to the developing world. By one such definition found in the literature barely one person in 20 is deemed to be middle class in the developing world, and that proportion has changed little between 1990 and 2005, though with population growth an extra 80 million people in the developing world joined this "Western middle class" over that period.

This paper has focused instead on a definition that seems more defensible in the context of developing countries. Someone is identified as being in the developing world's middle class if she lives in a household with consumption per capita between $2 and $13 a day at 2005 purchasing power parity. The lower bound is the median poverty line of developing countries while the upper bound is based on the US poverty line. In 1990, about one person in three in the developing world was middle class by this definition; by 2005 the proportion had risen to one-in-two. An extra 1.2 billion people joined the middle class over 1990-2005.

This expansion in the middle class was not simply a horizontal shift in distribution whereby all income levels rise by a similar proportion. Instead, the huge expansion in the size of the developing world's middle class reflects global distributional shifts that have entailed greater poverty reduction at low poverty lines than one would have expected under a distribution-neutral growth process. The modal income level has increased very little with the rise in the mean. Instead, fewer people live near the mode, and densities have become "bunched up" just above $2 a day.

The more dense middle of the distribution can thus be seen as an implication of a pattern of aggregate economic growth in the developing world as a whole that has favored very poor people. The high growth rates of China and India have played an important role in producing the middle-income bulge in the developing world as a whole; indeed, China alone accounts for half of the 1.2 billion new entrants to the middle-class over 1990-2005.

The vulnerability of this new middle class to aggregate economic contraction is obvious; one-in-six people in the developing world now live between $2 and $3 per day. As the developing world confronts the spillover effects of the global financial crisis, lower growth rates are expected, with forecasts being revised downwards rapidly (as I write in late 2008), though with positive growth still expected. Whether the overall pace of poverty reduction slows, or is

even reversed, will depend crucially on the incidence of the impacts of the crisis across countries. It is clear from the calculations reported in this paper that maintaining growth in Asia will be especially important to preventing a slowdown in overall progress against poverty. Continuing growth in China and India is still expected, albeit it at a lower rate.

An exploration of how distributions have changed across developing countries reveals that, although China and India have naturally carried a large weight in the aggregate outcomes, the expansion in the middle class is evident for about 70% of countries. Two types of countries are now found within the developing world, according to whether they have a large middle class or a small one. Over time, the developing world has become even more polarized between these two groups of countries.

One's vulnerable to lower growth depends of course on where one lives. The extent of the middle class within a country is relevant, but the nature of its relevance might not be as obvious as one thinks. There is a "middle-class advantage" in attaining pro-poor growth in that a given rate of growth tends to bring higher proportionate impacts on poverty in developing countries that started the growth process with a larger middle class. This middle-class advantage is mainly due to initial differences in the incidence of extreme poverty. Starting out with a high poverty rate makes it harder to achieve a given proportionate impact on the poverty count through economic growth. That is not, however, true for the <u>absolute</u> impacts, and in making cross-country comparisons of progress against poverty it is probably better to use absolute impacts. A given rate of economic growth tends to yield a larger reduction in the poverty rate in poorer countries. By the same token, poor people living in countries with smaller middle classes will be more exposed to aggregate economic contractions, not less.

References

Atkinson, Anthony B., 1987, "On the Measurement of Poverty," *Econometrica* 55: 749-764.

Banerjee, Abhijit and Esther Duflo, 2008, "What is Middle Class about the Middle Classes Around the World?", *Journal of Economic Perspectives* 22(2): 3-28.

Birdsall, Nancy, Carol Graham and Stefano Pettinato, 2000, "Stuck in the Tunnel: Is Globalization Muddling the Middle Class," Center on Social and Economic Dynamics, Working Paper 14, Brookings Institution, Washington DC.

Bourguignon, François and Christian Morrisson, 2002, "Inequality Among World Citizens: 1820-1992," *American Economic Review* 92(4): 727-744.

Bussolo, Maurizio, Rafale De Hoyos and Denis Medvedev, 2008, "Is the Developing World Catching Up? Global Convergence and National Rising Dispersion," Policy Research Working Paper 4733, World Bank.

Chen, Shaohua and Martin Ravallion, 2004, "How Have the World's Poorest Fared Since the Early 1980s?" *World Bank Research Observer*, 19/2: 141-70.

_____ and _____, 2008, "The Developing World is Poorer than we Thought, but no less Successful in the Fight against Poverty," Policy Research Working Paper 4703, World Bank.

_____ and _____, 2009, "China is Poorer than we Thought, but no Less Successful in the Fight Against Poverty," in Sudhir Anand, Paul Segal, and Joseph Stiglitz (ed), *Debates on the Measurement of Poverty*, Oxford University Press.

Easterly, William, 2001, "The Middle Class Consensus and Economic Development," *Journal of Economic Growth* 6(4): 317-335.

_____, 2009, "How the Millennium Development Goals are Unfair to Africa," *World Development* 37(1): 26-35.

Ferreira, Francisco and Martin Ravallion, 2009, "Poverty and Inequality: The Global Context", in Wiemer Salverda, Brian Nolan and Tim Smeeding (eds), *The Oxford Handbook of Economic Inequality*, Oxford: Oxford University Press.

Kakwani, Nanak, 1993, "Poverty and Economic Growth with Application to Côte D'Ivoire," *Review of Income and Wealth,* 39, 121-139.

Lopez, Humberto and Luis Serven, 2006, "A Normal Relationship? Poverty, Growth and Inequality," Policy Research Working Paper 3814, World Bank.

Milanovic, Branko, 2005, *Worlds Apart: Measuring International and Global Inequality*, Princeton: Princeton University Press.

Milanovic, Branko and Shlomo Yitzhaki, 2002, "Decomposing World Income Distribution: Does the World Have a Middle Class?" *Review of Income and Wealth* 48(2): 155-78.

Pressman, Steven, 2007, "The Decline of the Middle Class: An International Perspective," *Journal of Economic Issues* 41(1): 181-199.

Quah, Danny, 1996, "Twin Peaks: Growth and Convergence in Models of Distribution Dynamics," *Economic Journal* 106: 1045-55.

Ravallion, Martin, 1997, "Can High Inequality Developing Countries Escape Absolute Poverty?" *Economics Letters,* 56, 51-57.

_____, 2001, "Growth, Inequality and Poverty: Looking Beyond Averages," *World Development*, 29(11): 1803-1815.

_____, 2007, "Inequality is Bad for the Poor," in J. Micklewright and S. Jenkins (eds.), *Inequality and Poverty Re-Examined.* Oxford: Oxford University Press.

Ravallion, Martin, Shaohua Chen and Prem Sangraula, 2008, "Dollar a Day Revisited," Policy Research Working Paper 4620, Washington DC, World Bank.

Ravallion, Martin and Michael Lokshin, 2007, "Lasting Impacts of Indonesia's Financial Crisis." *Economic Development and Cultural Change* 56(1): 27-56.

Sala-i-Martin, Xavier, 2006, "The World Distribution of Income: Falling Poverty and… Convergence. Period," *Quarterly Journal of Economics*, 121(2): 351-397.

Shukla, Rajesh, 2008, "The Great Indian Middle Class," National Council of Applied Economic Research, New Delhi.

Sridharan, E., 2004, "The Growth and Sectoral Composition of India's Middle Class: Its Impact on the Politics of Economic Liberalization," *India Review* 3(4): 405-428.

Thurow, Lester, 1987, "A Surge in Inequality," *Scientific American* 256: 30-37.

Wolfson, Michael, 1994, "When Inequalities Diverge," *American Economic Review* 84(2): 353-358.

_____, 1997, "Divergent Inequalities: Theory and Empirical Results," *Review of Income and Wealth* 43(4): 401-421.

World Bank, 2007, *Global Economic Prospects 2007: Confronting Challenges of the Coming Globalization*, Oxford University Press for the World Bank.

_____, 2008, *Global Purchasing Power Parities and Real Expenditures. 2005 International Comparison Program*, Washington DC: World Bank.

Zhang, Xiaobo and Ravi Kanbur, 2001, "What Difference Do Polarization Measures Make? An Application to China," *Journal of Development Studies* 37(3): 85-98.

Figure 1: B and C have less poverty than A, but B has a larger "middle-income bulge"
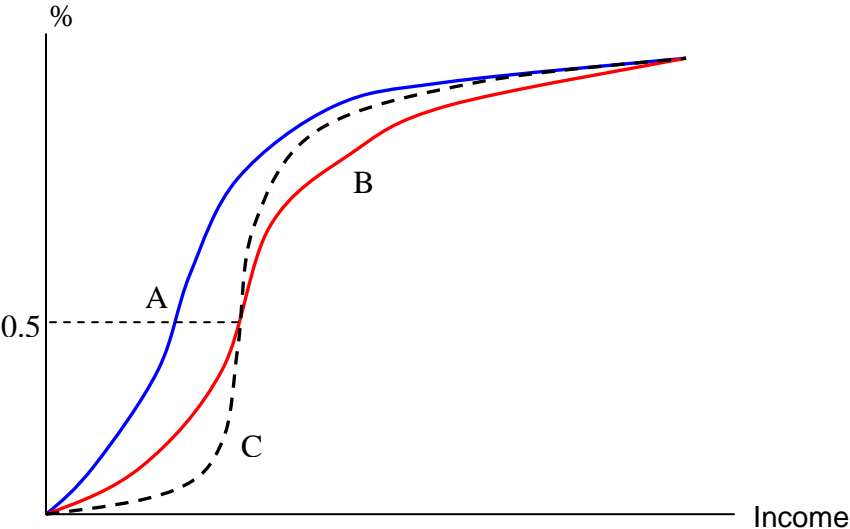
Figure 2: Distributions for the developing world as a whole in 1990 and 2005
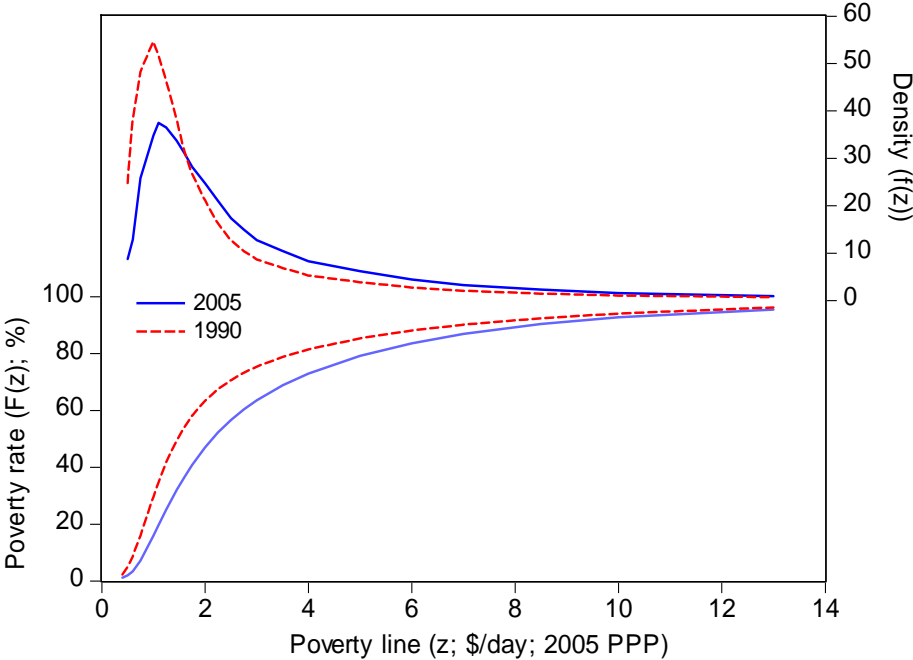
Figure 3: Assessing the bulge relative to distribution-neutral growth

Poverty rate in 2005 minus poverty rate in 1990 (% points)

Actual change

Simulated change under distribution-neutral growth

Poverty line ($/day; 2005 PPP)
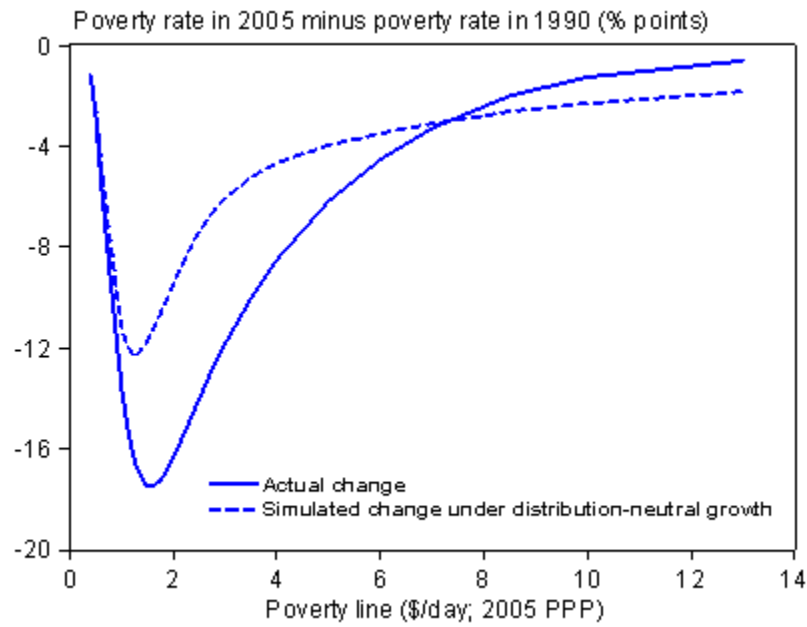
Figure 4: Elasticities of 2005 poverty rate to distribution-neutral growth

Elasticity (times -1)

Actual

Simulated under distribution-neutral growth

Poverty line ($/day; 2005 PPP)
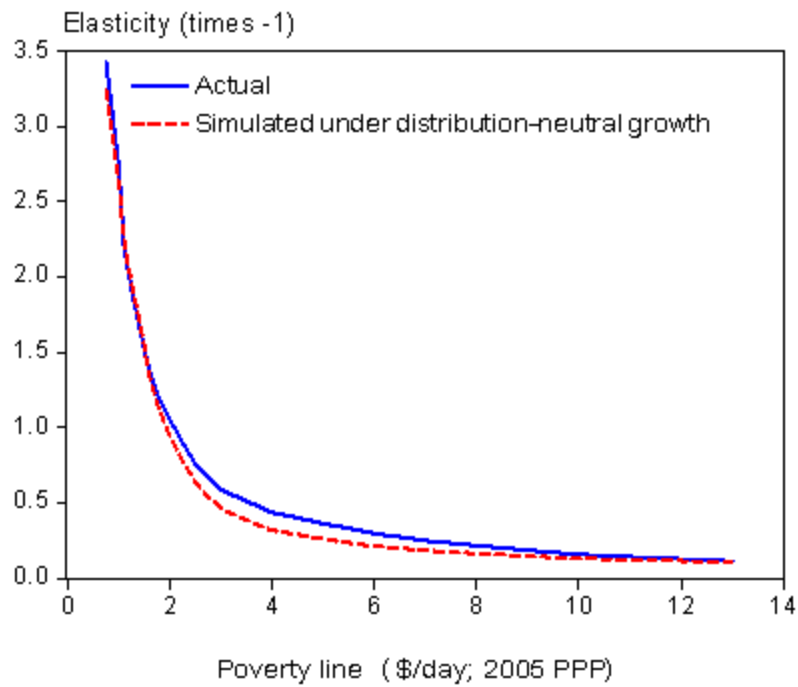
Figure 5: Middle-class share (*F*(13)-*F*(2)) across two surveys for same country

Middle-class share for latest survey (%)

— 45 degree line
— Least squares regression line

Middle-class share for earliest survey (%)
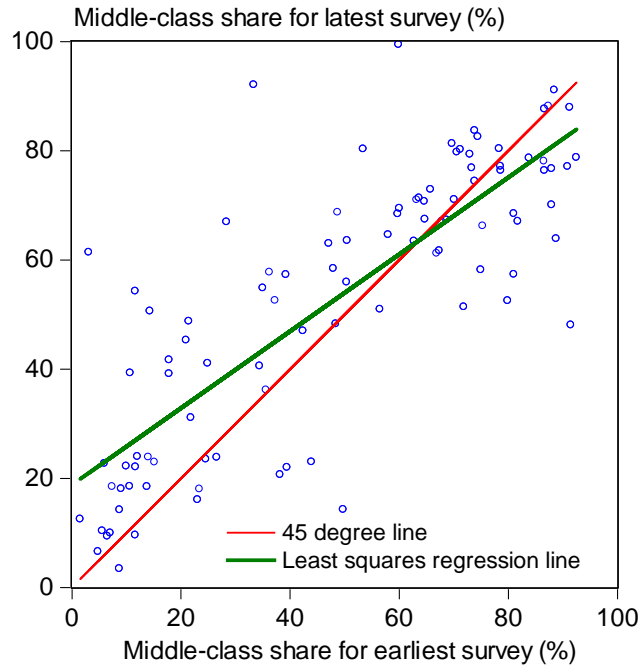
Figure 6: Densities (Epanechnikov kernel) of middle-class population shares

Kernel density

—— First survey's $2-$13 population share (%)
- - - Second survey's $2-$13 population share (%)
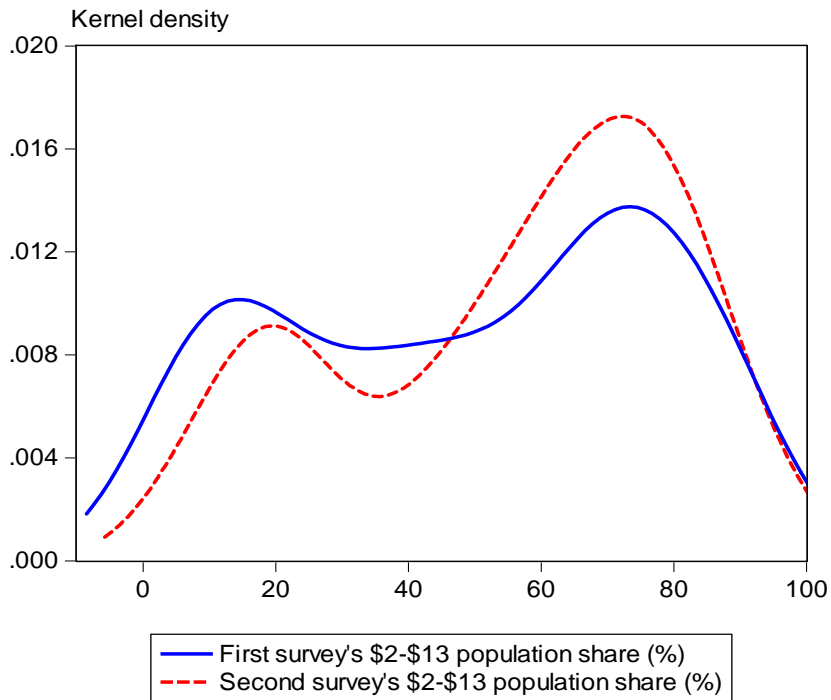
Figure 7: Greater middle income bulge in growing economies

Annualized proportionate changes in population
share between $2 and $13 per day

Annualized change in log survey mean
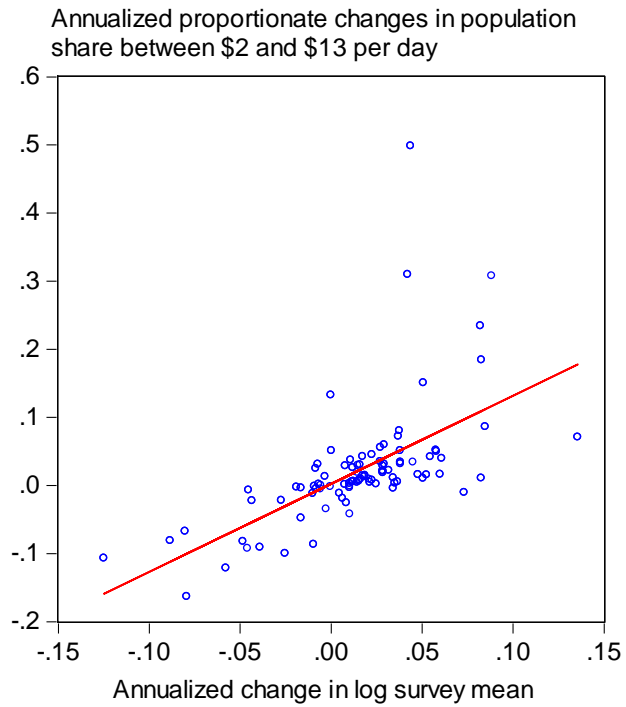
Figure 8: Changes in the Foster-Wolfson index of concentration against initial mean

Annualized change in log Foster-Wolfson index of
concentration around the median

Log household consumption per person in first survey

Table 1: Poor by the standards of developing countries

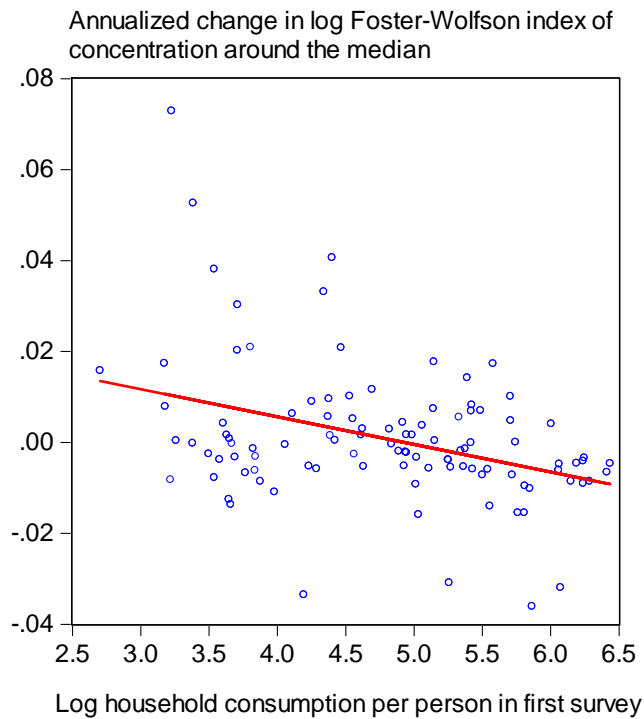| Region	Number living below $2 a day (millions)		Percentage of population	
	1990	2005	1990	2005
East Asia and Pacific	1273.7	728.7	79.8	38.7
Of which China	960.8	473.7	84.6	36.3
Eastern Europe and Central Asia	31.9	41.9	6.9	8.9
Latin America and Caribbean	95.9	94.3	21.9	17.1
Middle East and North Africa	44.4	51.5	19.7	16.9
South Asia	926.0	1091.5	82.7	73.9
Of which India	701.6	827.7	82.6	75.6
Sub-Saharan Africa	392.9	555.9	76.1	72.9
Total	2764.9	2563.9	63.4	47.0
Total excluding China	1804.1	2090.2	55.9	50.4

Source: Author's calculations using *PovcalNet*.

Table 2: Poor by US standards

Region	Number living below $13 a day (millions)		Percentage of population	
	1990	2005	1990	2005
East Asia and Pacific	1589.2	1845.8	99.6	98.0
Of which China	1134.5	1279.7	99.9	98.1
Eastern Europe and Central Asia	387.2	389.7	83.1	82.3
Latin America and Caribbean	372.6	456.5	85.1	82.9
Middle East and North Africa	214.7	291.6	95.2	95.5
South Asia	1118.6	1471.7	99.9	99.7
Of which India	848.3	1091.4	99.9	99.7
Sub-Saharan Africa	510.6	753.0	98.8	98.7
Total	4193.0	5208.2	96.1	95.5
Total excluding China	3058.5	3928.5	94.8	94.7

Source: Author's calculations using *PovcalNet*.

Table 3: Living between $2 a day and $13 a day

Region	Number living between $2 and $13 per day (millions) 1990	2005	Change 1990-2005 (millions)	Percent of total change	Percent of the population 1990	2005
East Asia and Pacific	315.5	1117.1	801.6	65.9	19.8	59.3
Of which China	173.7	806.0	632.3	52.0	15.3	61.8
Eastern Europe and Central Asia	355.3	347.8	-7.5	-0.6	76.3	73.4
Latin America and Caribbean	276.7	362.1	85.4	7.0	63.2	65.8
Middle East and North Africa	170.2	240.1	69.9	5.7	75.5	78.7
South Asia	192.7	380.2	187.5	15.4	17.2	25.8
Of which India	146.8	263.7	116.9	9.6	17.3	24.1
Sub-Saharan Africa	117.7	197.1	79.4	6.5	22.8	25.8
Total	1428.1	2644.3	1216.2	100.0	32.7	48.5
Total excluding China	1254.4	1838.3	583.9	48.0	38.9	44.3

Source: Author's calculations using *PovcalNet*.

Table 4: Joint distribution of countries by size of their middle-class

Count of countries		Middle-class population share in latest survey ($F_t(13) - F_t(2)$, %) [0, 20)	[20, 40)	[40, 60)	[60, 80)	[80, 100)	Total
	[0, 20)	12	8	3	1	0	24
Middle-class population	[20, 40)	2	6	8	1	1	18
share in earliest survey	[40, 60)	1	1	5	5	2	14
($F_{t-\tau}(13) - F_{t-\tau}(2)$, %)	[60, 80)	0	0	3	19	5	27
	[80, 100)	0	0	2	10	4	16
	Total	15	15	21	36	12	99

Source: Author's calculations; see text.

Table 5: Regressions for middle-income bulge over time

	Change in log ratio of proportion under $13 to proportion under $2 (1)	(2)	Change in log Foster-Wolfson index of concentration around the median (3)	(4)
	Survey mean	National accounts consumption	Survey mean	National accounts consumption
Intercept	-0.086*	-0.076*	0.031**	0.021*
	(-2.372)	(-1.982)	(2.995)	(2.095)
Growth rate (annualized change in log mean, $g_i(\mu_{it})$)	1.384**	0.830**	-0.028	-0.011
	(5.893)	(2.977)	(-0.698)	(-0.277)
Initial log mean ($\ln \mu_{it-\tau}$)	0.019*	0.017	-0.006**	-0.004*
	(2.327)	(1.892)	(-2.993)	(-2.085)
N	91	86	91	83
R^2	0.371	0.169	0.117	0.066

Notes: The "mean" (both initial value and growth rate) is the survey mean in regressions (1) and (3) and private consumption per capita from the NAS in (2) and (4). The t-ratios based on robust standard errors in parentheses. * denotes significant at the 5% level; ** denotes significant at the 1% level.
Source: Author's calculations; see text.

Table 6: Regressions for change in poverty rate as a function of growth rate and initial poverty level

| | Annualized change in poverty rate for $2 a day | | | | | |
| | Change in log poverty rate ($g_i(F_{it}(2))$) | | | | Change in level of poverty rate ($d_i(F_{it}(2))$) | |
	(1) OLS	(2) IVE	(3) OLS	(4) IVE	(5) OLS	(6) IVE
Intercept	-0.012	-0.005	-0.012**	-0.008	0.0201	0.003*
	(-1.908)	(0.607)	(-2.175)	(-1.365)	(1.863)	(2.407)
Growth rate (annualized change in log survey mean, $g_i(\mu_{it})$)	-2.615**	-3.323**	n.a.	n.a.	-0.207**	-0.242**
	(-6.608)	(-4.560)			(-4.203)	(-3.152)
Growth rate interacted with initial poverty rate ($g_i(\mu_{it}).F_{it-\tau}(2)$)	2.621**	3.101	n.a.	n.a.	-0.277**	-0.386**
	(4.915)	(3.746)			(-3.035)	(-2.904)
(1-Poverty rate) *times* growth rate ($g_i(\mu_{it}).(1-F_{it-\tau}(2))$)	n.a.	n.a.	-2.613	-3.294	n.a.	n.a.
			(-7.273)	(-4.585)		
N	91	86	91	86	99	91
R^2	0.535	0.458	0.535	0.466	0.725	0.535
Homogeneity test	0.037	-0.620	n.a.	n.a.	6.970	5.623

Notes: The t-ratios based on robust standard errors in parentheses. * denotes significant at the 5% level; ** denotes significant at the 1% level.
The homogeneity test is the t-test for the sum of the coefficients on the growth rate and the growth rate interacted with initial poverty rate.
Source: Author's calculations; see text.